IN THE UNITED STATES DISTRICT COURT
EASTERN DISTRICT OF VIRGINIA
NORFOLK DIVISION

```
CENTRIPETAL NETWORKS, INC.,   )
                              )
          Plaintiff,          )
v.                            ) Civil Action No.:
                              )     2:18cv94
CISCO SYSTEMS, INC.,          )
                              )
          Defendant.          )
```

TRANSCRIPT OF VIDEOCONFERENCE BENCH TRIAL PROCEEDINGS
(**Confidential Testimony Redacted**)

Norfolk, Virginia
May 12, 2020

Volume 5B
Pages 658-718

BEFORE: THE HONORABLE HENRY C. MORGAN, JR.
United States District Judge

Appearances: (Via Zoomgov Video)

KRAMER LEVIN NAFTALIS & FRANKEL, LLP
By: JAMES RUSSELL HANNAH
Counsel for Plaintiff

DUANE MORRIS, LLP
By: MATTHEW GAUDET
Counsel for Defendant

I N D E X

E X H I B I T S

PROCEEDINGS

(Proceedings resumed at 2:02 p.m. as follows:)

THE COURT: All right. Counsel, do you want to have a conference closed to audio?

Would you close the audio?

COURTROOM DEPUTY CLERK: Yes, sir.

(Confidential Testimony to Page 669, Line 5 redacted and filed under seal.)

* * *

* * *

BY MR. HANNAH:

Q. So Doctor, based on the evidence that you've seen, do the firewalls receive a first and a second rule set?

A. Yes. Both the firewalls themselves and the management centers will also be sort of continually updating or receiving rule set information.

Q. If we turn to the next element, which is the preprocessing of the first and second rule set to optimize the performance, do the firewalls with the Firepower Management Center meet this limitation?

A. Yes, they do.

Q. All right. If we could go to PTX-1289?

MR. HANNAH: And Your Honor, this has already been admitted into evidence.

THE COURT: Right.

BY MR. HANNAH:

Q. If we could go to Page 1594 which ends in the same corresponding Bates label, 1594.

MR. HANNAH: Your Honor, it's the diagram that we were talking about before the lunch break.

THE COURT: Right.

BY MR. HANNAH:

Q. Doctor, can you please explain where the preprocessing of the rules in order to optimize them occurs in this diagram?

A. Right. So that occurs where we see the Threat Intelligence Director. So it's again pulling in, getting what's labeled in this diagram as indicators, right, and sending out what are labeled as observables down to the devices. So observables are you know -- indicators can be more complex rules, but what the Threat Intelligence Director does is preprocess, simplify, optimize, choose, remove duplicates and so on, and correspondingly will pass observable information to the managed devices.

THE COURT: Managed device is the firewall, right?

THE WITNESS: Exactly. Yes. Sorry, I should have said that.

THE COURT: All right.

BY MR. HANNAH:

Q. So if we go to -- I'd like to turn your attention to PTX-1393, and we can take a closer look at the observables that you just talked about.

Can you please explain what 1393 is, Doctor?

A. Certainly. So 1393 refers to Lamplighter. Lamplighter, remember, as I said, is another name for the Threat Intelligence Director. And again, you can see from its description that it matches what we've been saying. Lamplighter is a system for consuming and normalizing heterogeneous third-party cyber threat intelligence, publishing the intelligence to detection technologies, and correlating the observations made by the detection technologies. And we'll see, you know, a bit more about what that means inside when it discusses observations.

MR. HANNAH: Your Honor, at this point I'd like to move PTX-1393 into evidence, please.

MR. GAUDET: No objection.

THE COURT: That will be admitted.

(Exhibit PTX-1393 received in evidence.)

MR. HANNAH: Thank you, Your Honor.

If we could go to Page 9 of this document? And Your Honor, it has the same corresponding Bates label, which is 009.

BY MR. HANNAH:

Q. If we look under 1393, and it's at Page 9, and if we look under observables and go all the way down to -- up to before the elements -- yes. Can you explain what's being shown here in terms of both the receiving of the rules that we just talked about and the preprocessing and optimization of the rules for the current element?

A. Right. So if you see, an observable is a basic unit that can be shipped to the elements, okay? So this is something that we're going to ship out to the firewalls. But the first thing that it has to do is it has to receive them, right? It has to receive them in a feed. And you can see that reads in these information, and you can see that the information associated with an observable includes an action, right, such as monitor or block, right? So this information that it can be getting or getting fed in as threat intelligence can include monitor or block information. So, actions associated with the observable. And you can see that it has to read in and process. That's the next step where it talks about the Json representation. I forgot now what Json stands for. Json is just a way of representing data in sort of, as you can see, a tabular type form. And it's read in -- after it's read into the feed it has a global ID, and in particular it tells you that this is an observable and it has an associated action with it.

Q. When it says the word ingester, is that what's being receiving or absorbing these, this information?

A. Yes. That's the part that receives or absorbs the information up in the Threat Intelligence Director.

Q. And is it your understanding that Json stands for Javascript Object Notation?

A. Javascript Java Notation. That's a way of representing data that's useful to many computer systems.

Q. And if we go further down in this document, it says "shipped to elements." What does that mean? What are the elements here?

A. Again, the elements would correspond to, correspond to the firewalls. You can see that what's shipped out is rule information.

Q. And we see that it has the "Action":Block there as well?

A. Yes.

Q. All right. Doctor, I'd like to move on to some testimony and?

MR. HANNAH: In your binder, Your Honor, it's PTX-1918. This is the testimony of Michael Sheck. It is from December 11th, 2019 and it ranges in Page 66, Line 13 through Page 67, Line 12. And Michael Sheck is the Cisco director of the Incident Response Team.

THE COURT: All right. That will be admitted.

MR. HANNAH: Thank you, Your Honor.

(Exhibit PTX-1918 received in evidence.)

BY MR. HANNAH:

Q. So before we have you explain, there's an acronym in here, Doctor. Does what does FMC stand for in this?

THE COURT: I looked that up. That's Firepower Management Center.

MR. HANNAH: Thank you. Yes, Your Honor.

BY MR. HANNAH:

Q. So Doctor, can you please let us know how this informed your opinion?

A. Right. So first he's asked what does FMC do to the threat intelligence feeds received from the Threat Intelligence Director. So you know, he discusses if you were to configure your FMC to receive a feed from the Threat Intelligence

Director. So first this just confirms what we've been saying. The Firepower Management Center receives feeds by way of, you know, the component that does that on the FMC is called the Threat Intelligence Director. Then he says you could use those threat indicators to fire a signature and detect traffic. Then he's asked further about that. But he says that signatures are essentially a set of logic contained on an intrusion detection or intrusion prevention system that look for specific sets of, of characteristics, and if there's a match to those characteristics, it creates an alert.

So here he's describing rules, right? That you use the FMC, that it takes those threat indicators and turns them into the appropriate rule format for the firewall, for the corresponding device. And I call this a rule because he says it's a set of logic corresponding to characteristics, and if there's a match it creates an alert. It takes an action. We've also seen besides creating an alert it might conceivably block. It depends on the setting. But here he's talking about an alert, but that's still a corresponding action and corresponds to a rule.

THE COURT: I thought the Threat Intelligence Director sent the intelligence directly to the firewall? I didn't realize it had to go -- which seems to me backwards -- through the FMC to the firewall.

THE WITNESS: Right. So the Threat Intelligence

Director is, as we've seen in the diagrams, is on the FMC. It's a part of the FMC. It's just given its own label or component name because of its very visible role. And yes, besides reading it in, it does have to do some processing and optimization.

THE COURT: Well, does that involve sending it to an intermediary before it goes to the firewall?

THE WITNESS: No. The Threat Intelligence Director is on the Firepower Management Center, and the preprocessing and optimization I'm talking about is done on the Firepower Management Center.

THE COURT: So it goes to the Firepower Management Center before it goes to the Threat Intelligence Detector?

THE WITNESS: No. So the TID is part of the Firepower Management Center. So the data feed comes in to the TID, but that's part of the Firepower Management Center. It's absorbed and adjusted, sort of the words that are used. It's then, all that information is -- for instance it's shown it's put in a database, it's recorded. Based on what comes in, there's further processing and optimization to determine what stuff should actually to out. What signatures are important to put on the device. And the reason you would do that is because, you know, it may be that you get the same intelligence from multiple sources, and in that case you wouldn't want to repeatedly send it out, for instance.

THE COURT: Okay.

BY MR. HANNAH:

Q. And maybe we can just go back to the demonstrative that we've shown for the Firepower which shows the threat intelligence coming in to the Firepower Management Center. Maybe that will help just illustrate.

Doctor, can you explain the flow in terms of how it works in this demonstrative?

A. Right. So we have the threat intelligence coming in, and it's shown as coming in to the Firepower Management Center, but remember, the Threat Intelligence Director is a part of that, is a component in that. And so the threat intelligence is coming in to the Threat Intelligence Director, processed throughout the Firepower Management Center, and then the appropriate information is then distributed out to the firewalls down below.

THE COURT: All right.

MR. HANNAH: All right, Doctor. I'd like to show you one more piece of source code for this element, and so at this point I'd like to seal the courtroom and mute the line, and our corporate representative Jonathan Rogers has left the room, or left his room.

If I may proceed, Your Honor?

THE COURT: All right.

(Confidential Testimony to Page 679, Line 1 redacted and filed under seal.)

* * *

* * *

BY MR. HANNAH:

Q. We turn back to the claims. Doctor, is it your opinion that the firewalls, including the ASA and the Firepower with the FMC, preprocess the first rule set and the second rule set to optimize performance?

A. Yes. That's what's being done on the FMC. It's preprocessing the rule sets as they come in to optimize their performance before sending them on to the firewall devices.

Q. Can we check that box?

A. Please.

Q. Thank you. Let's move on to the next set of elements. And again, we grouped these together. Can you just briefly remind the Court why we grouped these together?

A. Right. So all of these deal with the issue of just using the first rule set. And so here, you know, we're discussing the configuring of and the processing of packets using the first rule set.

Q. So if we turn to PTX-1293. PTX-1293, can you please explain what this is, what this document is, Doctor?

A. This is another configuration guide for the ASA appliances that we're discussing.

MR. HANNAH: Your Honor, at this point we'd like to move PTX-1293 that evidence.

MR. GAUDET: No objection.

THE COURT: All right. 1293 will be admitted.

MR. HANNAH: Thank you, Your Honor.

(Exhibit PTX-1293 received in evidence.)

MR. HANNAH: If we could turn to Page 668 of this document? Your Honor, it has the same corresponding Bates label of 0668.

BY MR. HANNAH:

Q. And there's a section that says choose a rule engine, transaction, transactional-commit model. Do you see that, Doctor?

A. Yes.

Q. Can you please explain to the Court what is this transactional-commit model and how this informed your opinion with regard to the claim elements of configuring and using the first rule set.

A. So this transactional-commit model is what it calls, I guess for these products, what they call the sort of the swapping rule. So the idea of the transactional-commit model is that, as you can see below where it says you can change this behavior, so that the rule engine uses a transactional model when implementing changes, continuing to use the old rules until the new rules are compiled and ready for use. With the transactional model, performance should not drop during the rule compilation. So the idea is that their previous method would sort of try and fold in the new rules along with the old rules, and then that could lead to these performance costs that it describes up above. So what they have moved to is this transactional-commit model where the idea is similar to what we talked about with the previous sorts of rules and rule sets, is that you get all the rules ready, you do a swap, and then you return back to the new rules completely.

Q. Is this transactional-commit model similarly used in that you can swap the rules without dropping any packets?

A. Yes. One of the problems that they were having previously was that their previous system would, in cases -- the performance cost would lead to the dropped packets.

Q. And if we go to the next page, which is 669 --

MR. HANNAH: And Your Honor, again, it has the corresponding Bates label of 0669.

BY MR. HANNAH:

Q. -- and looking at that first paragraph, how does this inform your opinion with regard to the transactional-commit model and the processing of a first rule set and being able not to drop packets during the swap?

A. So, "An additional benefit of the transactional model is that, when replacing an ACL -- that's Access Control List, or rules, on an interface -- there is no gap between deleting the old ACL and applying the new one. This feature reduces the chances that acceptable connections may be dropped during the operation."

So again, this shows that the old system that didn't use this technique would have the potential or be liable for dropping packets, which is not, as we've said before, something that you'd want, and here that there's sort of a more natural -- you get all the new rules ready to go, and then you swap them in and then you go from doing the old rules to the new rules directly, and of course as we'll see, during that time you have to cease operations and cache packets. But then you can have a relatively seamless swap between the first and second rule set.

THE COURT: Well, this seems slightly different than the last system, because the last system seems to apply the new rules earlier in the process than this system. This says it continues to use the old rules while the new rules are being protected instead of stopping the flow.

THE WITNESS: It's still going to need to stop the flow when it does the rules, it just waits until the last possible moment to do so. So it waits until everything is ready before it does the stop, does the swap, and then moves on.

THE COURT: This seems slightly different --

THE WITNESS: It's definitely --

THE COURT: -- than the technology we talked about with regard to switches and routers.

THE WITNESS: It is. And that's why we've had to handle them separately. They are two different technologies underlying it. Because the firewalls work a bit differently than the routers and switches.

THE COURT: Well, it sounds as if the routers and switces apply the new rules earlier in the process than the firewall technology.

THE WITNESS: I, I -- so I think with the routers and switches it, again, has to go through the same aspect of making things ready for the TCAM, right? So remember, for the routers and switches, we saw that it had to prepare things, prepare the rules to be put --

THE COURT: Yes, but it did that while they were -- I'm used to saying "cached". But it did that while they were being cached. Which means that during that period of time the old rules weren't being applied, no rules were being applied. Now of course that's a very short period of time, but it seems like this one they apply the old rules instead of --

THE WITNESS: I think --

THE COURT: I don't see anything about caching them in this.

THE WITNESS: Well, so I think that's just because we're not there in our discussion yet. So I will be talking about that. But in particular, that has to do with the -- like the issue of there being no gap between deleting the old ACL and applying the new one. What that means is they're not going to drop packets and they're not going to -- they're not going to drop packets and they're not going to allow packets through without any rules. So as you've pointed out before, while these machines are very fast, there's always some amount of time between the time it takes for them to decide, okay, I'm moving from one rule set to another, and during that time it will cache the packet. And in that sense it's the same.

THE COURT: Doesn't seem the same thus far. But let's keep going.

MR. HANNAH: All right. Thank you, Your Honor.

BY MR. HANNAH:

Q. So I'd like to turn your attention to PTX-1241.

Doctor, can you explain what this is?

A. This is, again, a firewall configuration guide. So it describes aspects of configuring to a user.

MR. HANNAH: Your Honor, at this point I'd like to move PTX-1241 into evidence, please.

MR. GAUDET: No objection.

THE COURT: 1241. All right. PTX-1241 is Configuring Access Rules. This will be admitted.

(Exhibit PTX-1241 received in evidence.)

MR. HANNAH: Thank you, Your Honor.

Could you go to Page 4 of this document?

And Your Honor, that ends in the Bates label 253.

BY MR. HANNAH:

Q. And if we look at the bottom for the transactional-commit model, can you explain for the Court how this informed your opinion with regard to this element of configuring and using the first rule set?

A. Right. So what this is saying is that it's describing the transactional-commit model, and what it's saying is what it's trying, I was trying to describe before, that when it's enabled, what you do is when you receive the new rules from the Firepower Management Center, you prepare them for the rule-matching engine in the firewall. And that's called the compilation. And then once that's done you do a swap, okay? So it says a rule update is applied after the rule compilation is completed. And the idea here is, as the judge has pointed out before, to minimize the time as much as you possibly can between when you're going from the old to new rule set so that you can cache and keep all the packets without running into problems.

THE COURT: Let's go back. You said something about commit the new rules.

THE WITNESS: Right. It's called transactional-commit. So the idea is that you've received the new rules, you've prepared them, and then this commit stage is when you do the swap. You're saying, okay, I'm ready to do the swap, I'm going to swap the rules. Similar to before, you have to validate, check and then commit and say, okay, I'm now committed, I'm now ready to use the second rule set.

THE COURT: But while you're doing that you're using the old rule set.

THE WITNESS: No, you have to stop in order to do the swap.

THE COURT: While you're preparing --

THE WITNESS: While you're preparing.

THE COURT: -- the new rules.

THE WITNESS: While you're preparing the rules you are using the old rule set. But when you actually need to do the swap, then you have to stop and make the change. And that's the commit.

THE COURT: And you're saying this is the same; it would end up with the old rules being --

THE WITNESS: Tossed out.

THE COURT: -- applied in the same, to the same number of packets as it would in the technology we discussed using switches and routers? Are you saying they're the just the same?

THE WITNESS: I don't want to say they are just the same, because they are different products and they have some slightly different underlying technologies. But definitely the idea or the manner is the same in the sense that you're doing a sort of complete swap and so you have to stop, do the swap, and then move on to using the new rule set. So that's why you're using the old ones up until the time you're ready. When you're ready you say, okay, I'm going to stop everything, switch things over, and then continue.

BY MR. HANNAH:

Q. And maybe it'll help, Doctor, if we pull it into the claim language, which requires configure at least two processors to process packets in accordance with the first rule set. And then it says after the preprocessing of the first rule set and the second rule set, configuring the at least two preprocessors to process packets in accordance with the first rule set you receive the packets, and then you process them in accordance with the first rule set. What's the first rule set that we're pointing to here?

A. So the first rule set would be the initial one that would come in. It's what would be referred to in the documents as the old rule set. So you're working with the old rule set, right, and you're still using that old rule set while you're doing, getting ready the second rule set.

Q. And then I know this is further down in the claims, but just for context, what's the second rule set that you're talking about that gets applied?

A. The second rule set in that document corresponds to what are called the new rule set. So you have to switch over to that new rule set once it's ready and you have to -- the signaling, say it's ready, make the swap, signal when you're done, and then move on to the second rule set.

MR. HANNAH: Your Honor, I'd like to show the Doctor some source code for this element, so we would like to seal the courtroom. We have confirmed that Jonathan Rogers has left.

We'd like to mute the line.

COURTROOM DEPUTY CLERK: Give me a minute.

THE COURT: All right.

COURTROOM DEPUTY CLERK: Wait, wait.

MR. HANNAH: I'd like to show you -- Doctor, I'd like to show you --

THE COURT: Wait.

COURTROOM DEPUTY CLERK: Okay. You're good.

THE COURT: All right. You may proceed.

MR. HANNAH: Thank you, Your Honor.

(Confidential Testimony to Page 693, Line 21 redacted and filed under seal.)

* * *

BY MR. HANNAH:

Q. Doctor, did you also take some, did you also test the product in your analysis?

A. Yes. I performed some basic tests.

Q. So I'd like to show you PTX-408. Can you explain for the

Court what is PTX-408?

A. So this looks like a slide talking about -- or sorry this is just the cover for some screenshots I took while testing.

MR. HANNAH: Your Honor, at this point I'd like to move PTX-408 into evidence.

MR. GAUDET: No objection.

THE COURT: What is this? This is some test you made using Cisco products?

THE WITNESS: Yes. So we had available some of the Cisco products, and I used it as a chance to examine the system and perform some basic tests.

THE COURT: Okay.

(Exhibit PTX-408 received in evidence.)

BY MR. HANNAH:

Q. So Doctor, if we turn to Page 197 of this --

MR. HANNAH: And it ends, Your Honor, in Bates label 822 --

BY MR. HANNAH:

Q. If we look at the top, there's a tab that says the Firepower Management Center. Do you see that, Doctor?

A. Yes.

Q. Can you explain to the Court what's being shown on this slide?

A. So this is when I was working with the Firepower Management Center. You can see it's under Policies. And what I would do is just take a policy and like change one of the rules so there would now be a new rule set that would have to be implemented. And then I wanted to try and just see what would happen as I tried to push traffic through the network as fast as I could through the system and see what would happen. And in particular if there would be dropped packets. And you know, the test was just because I was using the transactional-commit model, there shouldn't have been any dropped packets, or I'd expect to see no dropped packets.

Q. Looking at this slide, does this show that there is the configuration of -- this claim element, the configuration of the processors in accordance with the first rule set and the receiving of packets and the processing of those packets?

A. Yes. So in particular it was receiving -- and we'll see this more on, I think, the next slide. This was receiving packets under a first rule set, then I would make a change and it would have to deal with the second rule set.

Q. So if we go to the next page of this which is ending in Bates label 823, what does this show, Doctor?

A. So this shows that the second rule set was deployed. It says deployment to device successful. So that would have been the deployment of the second rule set.

Q. If we go to the next page of this document, and if we look at the bottom where it says the dropped count?

A. Yes.

Q. How does that inform your opinion as to --

A. Well, right. This is just as I was running from one to the other, you know, what happened during the switch, did I see any dropped packets. And you know, just the answer was no, I didn't see any dropped packets as I was coming through, and the system is set up that I was keeping track of packets coming in versus packets coming out on the other side, and it didn't see any dropped packets.

Q. And Doctor, just to be clear, this is what you would see as a user using the Firepower Management Center to manage the firewalls?

A. Yes.

Q. Doctor, based on all the evidence that you reviewed, and going back to the claims, is it your opinion that the firewalls meet the "configure of the processors with the first rule set receiving a plurality of packets and processing those packets" as shown in the "configure after processing and process" limitations in claims 9 and 17?

A. Yes, I do.

Q. Can we check those boxes?

A. Please do.

Q. So let's turn to the next set of elements. And we have here where you signal each of the processors to process packets in accordance with the second rule set, and you configure the processors to process the packets with the second rule set. Do the firewalls meet these limitations?

A. Yes.

Q. If we turn back to PTX-1293?

MR. HANNAH: Your Honor, this document has been admitted into evidence already.

THE COURT: Right.

BY MR. HANNAH:

Q. So if we go to Page 668 of this document, can you explain, looking the transactional-commit model, how this is relevant to the elements that we're talking about here of configuring the processor to process packets in accordance with a second rule set?

A. Right. So again, we need to signal when everything is ready and then configure for the second rule set. And I think what I would point to is that where it talks about, you know, you can change this behavior, so that the rule engine uses a transactional model, right? So transactional, again, the implication is it's breaking things down into these phased steps when implementing rule changes, so it's going to use the old rules until the new rules are compiled and ready for use. So the ready for -- compiled and ready for use, like ready for use is the signal, right? So it's a signal that says, okay, you know, now these are ready, now let's go ahead and make the swap. So it's a signal that it's ready to do the swapping. And the fact that they're compiled means they're ready to go, so you can take the compiled rules and put them directly into the system so that the processors are ready to use those rules.

Q. So --

MR. HANNAH: May I move on, Your Honor?

THE COURT: Yes.

MR. HANNAH: Okay.

BY MR. HANNAH:

Q. So if we go to PTX-1196? And Doctor, can you please explain what this document is?

A. So this is tmatch, again discusses how, I guess, the transactional-commit model. You can see tmatch, transactional-commit model, software functional and design specification. So this is discussing the transactional-commit model that we've been discussing so far.

Q. And underneath that it says "An enhancement to avoid potential packet drop and reduce compilation time in some customer scenarios during large compilation of rules." Do you see that?

A. Yes. So in particular, as I think we've mentioned -- or we saw actually in the comments in the source code, a reason that they introduced this was because they were finding packet drops in certain situations which was causing some concern, I think.

MR. HANNAH: Your Honor, at this point we'd like to move PTX-1196 into evidence.

MR. GAUDET: No objection.

THE COURT: All right.

So they're trying to speed up the commitment process? Is this what this enhancement is for?

THE WITNESS: Yeah. So -- or what would happen I think, as what the other documents have said, is before what they would try and do is they would try and say, we'll have the new rules go as soon as we get them, but we won't actually, you know, do this preparation stage. And so they have to try and do them sort of together at the same time. They would say we're using the new rules even though they're not really ready, and we'll finish getting them ready in the background and start flushing out the old rules in the background. And it would create -- it created this big, I think, complicated mess. And the outcome of that complicated mess was that it would lead in situations to packet drops, which are exactly what you'd like, I think, to avoid.

THE COURT: Well, is the reason it led to packet drops because the old rules were still being applied? Is that a factor in increasing the packet drop?

THE WITNESS: I think at the end of the day the new rule would be applied, but it would just take so long that the packet would actually get dropped and the system would get overwhelmed. And so packet -- it wouldn't be able to run through the packets fast enough trying to do this hodgepodge of fixing when you were going to use the old rules, and making sure you were going use to use the new rules slowed everything down so much, even for that short period of time, that packets would drop.

THE COURT: Seems like we're sort of -- I'm a little mixed up here.

THE WITNESS: It's a complicated --

THE COURT: The answer to speed up the commit process, which sounds like it's a cousin of the old system. Because they're putting the rules in before they're quite ready if they're speeding up the commit process.

THE WITNESS: So I think part of it is that the way the speed comes about, or the way the effectiveness comes about is by the clear breaking things up into these fixed stages, including the signaling, the switching and the moving to the second. And before, they hadn't organized it in this way. They had sort of mushed everything together in a way that was much more complicated and less functional. And that led to the packet drops.

And I agree, it's really hard to explain -- like I would have a great deal of trouble trying to explain their old system to you, because their old system is very confusing. Like how do you stick in the new rules while the old rules are still there, and then if you somehow get, look like you're matching the old rule, make sure that, no, you're actually trying to match the new rule while you're doing all this other stuff of trying to get the stuff in the right order in the right sense. The old system was just trying to do 10 things at once, maybe, right? And that's why this breaking things down into these stages, like that's why it's called transactional-commit by Cisco. But it, the transactional-type process of step-by-step-by-step matches the claim language, and that's what I'm trying to show here.

THE COURT: All right.

BY MR. HANNAH: And so at this point Your Honor, I'd like to move PTX-1196 into evidence.

THE COURT: That will be admitted. I think I already did, but...

(Exhibit PTX-1196 received in evidence.)

MR. HANNAH: I'm sorry, Your Honor.

BY MR. HANNAH:

Q. So if we can go to Page 8 of this document. And I think, Doctor, so if we look at the top paragraph, can you explain how this informs your opinion with regard to the signaling of the processors to process packets with the second rule set and then configure those process packets in accordance with the rules -- the second rule set?

A. So it says with the legacy model, with the old model, new rules would take effect immediately during compilation. So this is what I was trying to say: They're trying to do 10 things at once. They're trying to say, okay, we'll get these new rules in even though we're still structuring them or trying to get them in the right format, you know, for the end device. Now because of that, the next sentence really describes the consequences. "However, during compilation, connections-per-second limit will be lower than the limit after compilation is done and the ASA becomes stable."

So what this means, again, this is under the old system, during this time the system's kind of unstable. It has to slow down. You know, this trying to do 10 things at once. I know this happens to me: If I try and do too many things at once, you know, things become unstable or I have to slow down. I can't really do them at the normal speed because otherwise things break. And in particular, the part that would, I guess, break under their system is that they would start to drop packets.

Now if we look at the next sentence "In contrast with the proposed transactional-commit model, new rules will not take effect until compilation is done and stable." So this gets back to the claim in the claim elements where it's like we're going to break things up into these pieces, we're going to wait until the rule set is ready, when we can signal that it's ready then we'll configure the second rule set, and we'll let everyone know the second rule set is ready, and then we'll move on to using the second rule set.

So that transactional, that step-by-step process that they've switched to, again, in my opinion, clearly aligns or matches the claim elements.

Q. So in the new system, Doctor, the transactional-commit model will signal to the system, which is this element, that there's a second rule set ready to go; is that right?

A. Right. It has --

MR. GAUDET: Your Honor, again, I think that one was leading.

MR. HANNAH: I can rephrase, Your Honor.

BY MR. HANNAH:

Q. So in terms of this line right here when it says the rules will not take effect until compilation is done and stable, can you explain how that meets the limitation of the "signal the processor to process packets with a second rule set"?

A. So right. That's obviously an English description and one I think that we can understand. But what it's saying is that you don't put in the new rules until the old rules are a compilation, everything is done and it's stable; that is, it's ready to go. When things are done and ready to go, that means you can signal, that means you give a signal saying, all right, we're ready to do the swap now. So this is the English description of what you would see or what corresponds to in the product as a signal that the second code, second rule set is ready to go.

Q. So let's take a look back at the claims.

Doctor, based on the documents that you've seen and the evidence in this case, do the firewalls signal each of the processors to process packets in accordance with the second rule set and configure the processors to process packets according to that second rule set?

A. Exactly. That's what -- you do the compilation, then as soon as the compilation is done, you give the signal, you say it's ready to go, we're ready to start moving on the second rule set, and as part of that you have to take that compiled second rule set, that rule set that's ready to go, and configure the processors to switch and use them, and flush out or remove the old rule set.

Q. Can we check those boxes?

A. Yes.

Q. I'd like to turn your attention back to PTX-1196 and go to the next page we were looking at, which is now Page 7 of the document. And for this -- I'm sorry, let's go back to the claims. Getting ahead of myself here.

So if we look at the next elements, we have "cease processing of the one or more packets and cache the one or more packets." Do you see that, Doctor?

A. Yes.

Q. Can you explain how the firewalls meet those two elements?

A. So as before, in some sense, when you're doing a switch, when you're doing this swap, that is going to take some action.

You have to deal with the memory holding the rules, remove the old rules, put in the new rules, verify that the new rules are ready to take effect, and then say that they're ready to go. During that time you can't be processing packets, right? Because you're actually trying to move, swap things in and out of memory. And that's, again, different than the legacy system. So during that time you ceased processing, and because it's not dropping packets, it has to keep those packets that are coming in that are waiting for processing or that are already there waiting for processing around until it's ready.

Q. So if we turn back to PTX-1196 and go to Page 7 --

MR. HANNAH: And as I stated, Your Honor, this is already in evidence --

BY MR. HANNAH:

Q. If you look at the -- let's look at the last line under the Customer Needs where it describes that the new system is -- "It is acceptable to have reasonable amount of delay for the new rules to take effect." Do you see that?

A. Yes. So what this is saying, if you look at the whole sentence, I mean, as long as the whole update transaction is atomic, no gap between removal of old rules and addition of new rules, it's acceptable to have reasonable delay for the new rules to take effect. That's matching what we're saying. The new rules come in, you compile and get them ready, and the idea is you have to swap out the old rules and the new rules.

There's not going to be a time -- in the old system you would have this drop potential where things could get dropped because things were not keeping up. And by doing this, and the word it uses here is no-gap, what you're doing is saying that I'm not going to have a period where I'm going to drop the packets.

I would like to be clear though that this no-gap is not saying that it happens instantaneously. Nothing even on a computer, which things happen very fast, happens instantaneously. When you're -- what it's saying is you're removing the old rules and adding the new rules. So you have to take the step, remove them out, put in the new rules, verify they have been added successfully, and then you signal that you're ready, and during that time, of course, the packets have to wait.

Q. And when the document talks about this reasonable amount of delay, what's it talking about there?

A. Again, this would be from the human standpoint, a very short amount of time, but from the computer standpoint it's a potentially significant amount of time where it actually has to wait for the new rules to take effect.

Q. And during that time it will cease processing packets, Your Honor -- I mean is that right, Doctor?

A. During the time of the swap it will have to cease packets, because it's taking the old rules and it's putting the new rules in. There's nothing else it can do but to wait during that time.

Q. Doctor, I'd like to turn your attention to PTX-1277, which has already been admitted into evidence. If we look at Page 7 that we looked at earlier, can you explain how this informs your opinion with regard to cease processing of the packets and caching the packets?

A. Yeah. So in particular you can see that there's various memories associated with the processors up top that are holding the packets. So you have Random Access Memory associated with the CPUs, typically some sort of Random Access Memory or additional memory that's connected to the CPUs. That's what stores packets before while they're waiting for service.

THE COURT: CTU?

THE WITNESS: CPU. It's Central Processing Unit. So that's the two or more processors.

THE COURT: Oh, you were saying CPU?

THE WITNESS: Sorry, yes. CPU. It's connected to the Central Processing Units. Those are the processors that are returning the rules, and the RAM is holding the packets while it's waiting for the system to be ready to process them.

THE COURT: That's like the cache?

THE WITNESS: Yes. That would be the cache. It's a memory. That's right.

BY MR. HANNAH:

Q. So Doctor I'd like to --

MR. HANNAH: Do you have another question, Your Honor? I'm sorry. Okay.

BY MR. HANNAH:

Q. If you turn to Page 12 of this document, if we look at the top portion it says "Packets arrive on the network interface card (NIC) are placed by the hardware on to a receive, Rxring." Do you see that?

A. Yes.

Q. So can you tell us, how did that inform your opinion as to whether the packets are cached?

A. So that's saying packets arrive and they're placed by the hardware. This receive ring, that's the cache. So sometimes caches are organized in certain ways. The typical way they're organized in this sort of system is the memory ends up looking like a ring. So it's called a ring. And in particular it's called a receive ring because that's what's receiving packets. The receive ring is the, you could think of as the cache or the buffer that we're talking about.

Q. If we can turn back to the claims then?

Doctor, based on the documents and evidence and your testing of the products, do the firewalls meet the "cease processing of the one or more packets and caching the one or more packets" limitations for both claims 19 and 17?

A. For 9 and 17, yes, they do.

Q. So we turn to the last group of elements. Can you explain whether these elements are met, and just briefly to remind the Court, why you grouped these together for these last elements?

A. Sure. So much as before, the first set of rules was, we grouped together the elements all relating to the use of the first set of rules. Here we're doing the same with the second set of rules, okay? So you have to make sure everything's set up to process packets with the second rule set and let everyone know that the second rule set is ready to go. That's the signal completion of the reconfiguration. And then once that's done, start processing packets in accordance with the second rule set.

Q. If we turn back to PTX-1293, which is already in evidence, and if we go back to Page 669, which ends in Bates number 669, the same corresponding Bates label, and we look at that paragraph, the top full paragraph, at what point in the process are we talking about here when the second rule set is applied?

A. So again, it's described in the chart roughly speaking as after compilation. So after compilation you go through that process, you signal the new rules are ready, the rules get installed, and then you start applying the new one once it's ready. And that's described in the sentence where it says that "an additional benefit of the transactional model is that, when replacing a ACL on an interface, there's no gap in deleting an old ACL and applying the new one." So as soon as you can get that signal that the second one's ready, you're back up and running again.

Q. And so we're clear, the second rule set would be the new rule set that's being applied?

A. The second one, sorry, is the new rule set, yes.

Q. If we go back to 1241 --

MR. HANNAH: Which has already been admitted, Your Honor.

BY MR. HANNAH:

Q. -- if we turn to Pages 4 and then over to 5, we didn't see Page 5 earlier, if we look at where it talks about the transactional-commit model at the bottom?

THE COURT: What's that Bates number?

MR. HANNAH: It is pages ending in Bates label 253 and 254. Bates label 253 and 254.

THE COURT: All right.

BY MR. HANNAH:

Q. And Doctor, can you explain how this passage is relevant to your opinion that the products are going to reconfigure to process packets in accordance with the second rule set, signal completion of reconfiguring to process the packets with the second rule set, and then begin to process the packets with the second rule set?

A. Here we can see in the first two sentences it's talking about the transactional-commit model. And as it says "when this feature is enabled, a rule update is applied after the rule compilation is completed."

So what we're talking about here, the reconfiguration and signaling that it's ready, that's the rule update, right? A rule update is what gets rid of the old rules, makes sure the new rules are up and ready to go, and it refers to that as the rule update.

Q. And then also can you take a look at the line that says "preventing packet drops while compiling large rules during high traffic rates", how did that inform your opinion?

A. Again this is saying unfortunately that the old system, as we've seen in the discussion, they noticed with customers that could lead to certain packet drops, and this system was designed to not have those packet drops by breaking it into the commit model framework.

Q. All right. If we go back to the claims, Doctor, based on all the evidence that you've seen, does the firewall -- do they reconfigure -- satisfy the reconfigure the signal completion and then responsive to receiving the elements of both claims 9 and 17?

A. Yes, they do.

Q. So could we check those boxes?

A. Please do. All three.

Q. So I'd like to just do a recap of your opinion with regard to the firewalls and why they meet the claims of the '806 patent. So could you explain what's being shown on this slide?

A. Certainly. So here we're dealing with the firewalls, and of course the firewalls in conjunction with the Firepower Management Centers are systems that provide network security. As systems they also have computer-readable media containing instructions that tell them what to do.

The Threat Intelligence Director receives rule sets from various sources, and it prepares them, preprocesses and optimizes them for distribution to the firewalls that it manages.

The firewalls will be using rule sets, so they will have a first rule set that they're using until a new rule set becomes available to them or they're told to update from the Firepower Management Center. When this new rule set is available, they use this transaction-commit model. So they say, okay, we're going to get the new rules, we're going to compile them; that is, put them into the right form for the corresponding memory where they will be stored, and then we'll do a swap of the rule sets designed so that they don't drop any packets. And the reason that they don't drop any packets during the time period they're doing the swap is because they're using this packet buffer cache. We've seen it referred to as a, for instance, the receive ring. After the rule set is swapped, then it can be signaled everything's ready to go, and the firewalls can begin processing the packets again according to this new rule set.

Q. Thank you, Doctor.

So in addition to your opinion with regard to literal infringement, did you provide an opinion with regard to the doctrine of equivalents?

A. Yes, I did.

Q. Can you please tell us, what is that opinion?

A. Yeah. So it might help if we just briefly put up the claim slide so I can point to the right element?

MR. HANNAH: The first one?

THE WITNESS: The first one, yeah. Previous.

A. Right. So this has to do with the preprocessing the first rule set and the second rule set to optimize performance of the commuting system for processing packets in accordance with at least one of the first rule set or the second rule set. So I guess now if we turn back, so I believe that I've shown this literally -- or I believe that there's literal infringement. I'm aware that the other side may be contesting that claim element with regard to their saying that these are not actually rules or that they're not being preprocessed and optimized. I disagree. Again, I think I've shown that and that it's a literal infringement, but if not, I would certainly want to put in an argument under the doctrine of equivalents.

So this works in substantially the same function; that is, it's receiving information regarding threats and threat information, including actions to be taken such as blocking and monitoring such bad traffic or potential threats. So the goal here is substantially the same function; that is, to turn this information into rules or to take these rules that it's receiving and reprocess them into a format appropriate for firewall.

BY MR. HANNAH:

Q. And Doctor, does it perform in substantially the same way?

A. Yeah. Substantially the same way. It ingests, absorbs the threat information, it preprocesses them, turns them into a form such as observables that are meant to be optimized for the system that they're going to run on for the corresponding firewalls. So it does this the same way as suggested by the claim element; that is, it's preprocessing this information, and further it's optimizing this information.

And the result is the same: That a rule set is passed down into the main firewall system where that system then puts that rule into use for the rest of the claim elements that follow. In particular, that could be the first rule set or the second rule set that are put into the system depending on the appropriate time.

Q. All right. Thank you, Doctor.

MR. HANNAH: With that, Your Honor, we'd like to turn to the last patent, which is the '205 patent.

THE COURT: Well, we're going to turn to adjourning.

MR. HANNAH: Okay.

THE COURT: I'll ask counsel to remain; however, the hearing will be adjourned and we'll turn off the audio.

COURTROOM DEPUTY CLERK: Audio is gone.

THE COURT: The amount of time it's taken for these two patents is unacceptable. I did not want to impose time limits because of the nature of the proceeding. This is the fifth day of the hearing. I don't know there's any diminution in my ability to understand the witnesses or observe them for purposes of assessing the weight to attach to their testimony. But I'm going to instruct counsel to present the Court with a schedule tomorrow morning as to when the plaintiff anticipates completing their case. We've only completed two out of five patents. I say to counsel who are used to practicing in this court that your chances of prevailing vary in inverse proportion to the weight of your filings in grams. I could say the same thing for the weight of your exhibits and the weight of the transcripts of the testimony in grams. I understand your evidence, but it's extraordinarily repetitive, and I think it can be presented in a far more efficient manner.

The protocols that were mostly set up by counsel I think have worked very well. Everything is as organized as far as my being able to put my hands on exhibits, it's as good as or better than most cases I try with everybody seated in the courtroom. It's gone, if anything, better than usual. But at the rate we're going, this case will take five weeks at least instead of three. And there's no reason for that. I mean, we haven't even heard one word of cross-examination of this witness after all the time he's been testifying, and he's only covered two of the three patents. We cannot continue to operate at this pace. So counsel will have ready for me tomorrow morning a schedule of their remaining witnesses and how long it's going to take to complete those witnesses. Because we cannot continue at this pace.

I'll ask, after counsel for the plaintiff submits it and we have a chance to discuss it tomorrow, we'll talk about counsel for the defendant. But I'll have to get it from the plaintiff at the beginning of court tomorrow. It's just not necessary to go into as much time as has been consumed presenting these two patents. It's just way too detailed. And you've got an intelligent witness who understands the process. And it just can't continue to take this long.

Do you have anything you'd like to say about that?

MR. HANNAH: Your Honor, we'll have the schedule ready for you in the morning.

THE COURT: All right. Anything from the defendants?

THE DEFENDANT: No, Your Honor. Thank you.

THE COURT: All right. I know it's often customary to set times in patent cases, and I've done that sometimes. Not always. And this is a complex case. But it's not this complex. And I don't think we're losing anything at all, as I say, by proceeding electronically as opposed to having everybody here in the courtroom. I think it's worked very efficiently. And I'm sure we'll be seeing a lot more of this. I mean, when you think of the witnesses we've had from all over the country, the lawyers we've had from all over the country, it's really gone very efficiently, I think. And of course part of that are the protocols that counsel worked very hard to put together. So... But we've just got to do better time-wise.

So we'll be adjourned until tomorrow morning at 10.

(Whereupon, proceedings concluded at 3:49 p.m.)

CERTIFICATION

I certify that the foregoing is a true and correct redacted transcript of Volume 5B of the proceedings held in the above-entitled matter.

Paul L. McManus, OCR

Digitally signed by Paul L. McManus, OCR
DN: cn=Paul L. McManus, OCR, c=US,
email=pmcmanusocer@verizon.net
Date: 2020.08.10 10:28:37 -04'00'

Paul L. McManus, RMR, FCRR

Date